SUPERCOMPUTER

Contribution

Amdahl's Law and the statistical content of the NAS parallel benchmarks

Horst D. Simon

Applied Research Branch
Numerical Aerodynamic Simulation (NAS)
Systems Division, NASA
Ames Research Center
Mail Stop T27A-1, USA

Erich Strohmaier

Computing Center,
University of Mannheim,
D-68131 Mannheim,
Germany

The NAS Parallel Benchmarks have been developed at the NASA Ames Research Center. In the last three years extensive performance data have been reported for parallel machines both based on the NAS Parallel Benchmarks [1, 2] and on LINPACK [3]. In this study we have used the reported benchmark results and performed a number of statistical experiments. These included cluster, factor and regression analyses. We did this to find out how many of the NAS Parallel Benchmarks are – in a statistical sense – necessary, to represent all the reported results. We also fitted Amdahl's Law to the data, to see whether it is meaningful to apply more sophisticated performance models to the reported results. All statistical experiments were done for absolute performances as well as for the corresponding efficiencies. The analysis of Amdahl's Law was performed for both classes (Class A and B) of the NAS Parallel Benchmarks.

As parallel systems became more and more wide spread within the last years the interest in benchmark data of parallel systems increased. One of the best known and commonly used benchmarks in this area is the set of the NAS Parallel Benchmarks [1, 2]. This set of 8 "paper and pencil" benchmark problems has been developed at the NASA Ames Research Center. The latest results are available electronically on the WWW at the URL address: `http://www.nas.nasa.gov/NAS/NPB/`.
Another very common benchmark is the LINPACK benchmark [3], which has been in use for more than ten years. Results are also available electronically at the URL address:
`http://www.netlib.org/benchmark/to-get-lp-benchmark`.
As performing complex benchmarks on a parallel system can be very time-consuming for the implementor, one might wonder how many of the 8 NAS PBs are necessary to describe and represent the data and the characteristics of the different systems and how many of them can be explained by the results of the others. In this study we try to find out whether it is possible to reduce the number of benchmarks without losing information, and which benchmarks are similar. We did this by factor analyses based on the correlation matrix between the benchmark

Horst Simon is an employee of Computer Sciences Corporation. This work is supported through NASA Contract NAS 2-12961.

3
Znaty, E., *General presentation of the BECAUSE Benchmark Set: BBS*, Future Generation Computer Systems **10**, 365–379, 1994.

4
Dongarra, J.J., *Performance of* \
puters Using Standard Linear Eq
ware, Computer Science Depa
versity of Tennessee, CS-89-85,

5
Hockney, R.W. and C.R. Jesshope, *Parallel Computers 2*, Adam Hilger, Bristol, 1988.

6
Schönauer, W., *Scientific Compu*
tor Computers, North-Holland,
1987.

7
Schönauer, W., H. Häfner, L. Gross, M. Schmauder and R. Weiss, *Parallelization of black box software for elliptic and parabolic PDEs*, submitted for publication.

8
Hockney, R.W., *Computational S*
appear in Concurrency: Practice
ence.

9
Schönauer, W. and H. Häfner, *Explaining the gap between theoretical peak performance and real performance for supercomputer architectures*, Scientific Programming **3**, 157–168, 1994.

	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP
r_{max}	0.99										
n_{max}	0.87	0.82									
$n_{1/2}$	0.75	0.71	0.76								
EP	0.97	0.97	0.83	0.82							
MG	0.91	0.93	0.77	0.65	0.91						
CG	0.57	0.55	0.41	0.18	0.57	0.69					
FT	0.85	0.90	0.53	0.64	0.87	0.95	0.77				
IS	0.65	0.65	0.52	0.25	0.64	0.78	0.99	0.84			
LU	0.94	0.94	0.90	0.65	0.90	0.99	0.75	0.96	0.84		
SP	0.96	0.96	0.91	0.76	0.95	0.99	0.70	0.96	0.79	0.99	
BT	0.82	0.99	0.94	0.78	0.99	0.98	0.63	0.95	0.72	0.98	0.99

Figure 2. The correlation matrix of the NAS PBs of class B and the LINPACK benchmark results.

also used the correlation matrix between the efficiencies of the benchmarks with respect to the peak performance to eliminate the strong effect of the overall correlation to peak performance (Figure 3 and Figure 4). On the average the individual correlations are now smaller but in general still high. The peak performance shows no big correlations (neither positive nor negative!) to any of the efficiencies of the benchmarks. The correlations between the parameter $n_{1/2}$ and the efficiencies are now negative and stronger compared to previous cases but still not as high as the benchmark correlations. This means that no general conclusions about the efficiencies of benchmarks can be made from the LINPACK parameters n_{max} and $n_{1/2}$.

	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP
r_{max}	-0.06										
n_{max}	0.87	-0.17									
$n_{1/2}$	0.60	-0.55	0.66								
EP	-0.14	0.46	-0.15	-0.29							
MG	-0.18	0.78	-0.22	-0.46	0.61						
CG	-0.22	0.70	-0.32	-0.59	0.42	0.78					
FT	-0.16	0.74	-0.21	-0.51	0.57	0.89	0.88				
IS	-0.25	0.79	-0.27	-0.51	0.55	0.85	0.93	0.93			
LU	-0.22	0.79	-0.21	-0.58	0.29	0.86	0.86	0.83	0.84		
SP	-0.19	0.78	-0.18	-0.61	0.35	0.87	0.91	0.88	0.91	0.95	
BT	0.06	0.81	-0.07	-0.38	0.44	0.85	0.61	0.73	0.77	0.83	0.75

Figure 3. The correlation matrix of the efficiencies of the NAS PBs of class A and the LINPACK benchmark results.

The factor analyses of the benchmarks

Factor analyses can be used as an explorative method to get an overview on the structure of the given data, but cannot be used for testing or proofing any hypothesis. Therefore much care must be taken in interpreting the results.

results of the class A problem size. The result is checked by looking at linear regressions between different NAS PBs.

Amdahl's Law [4] gives a very simple model for the performance of a parallel system for different numbers of processors. We fitted the measured data to Amdahl's Law to see whether this is possible and whether sufficient statistical space for including additional parameters remains in this model. All analyses were done using the SAS statistical software package. The data used for the analyses in this paper is as of October 1994

The correlation matrices

As starting point for the factor analyses we had to calculate the correlation matrix. We used the NAS PB results of the class A benchmarks and the LINPACK results r_{max} from table 3 of the LINPACK report [3] for unlimited problem sizes. We included also the peak performance r_{peak}, and the parameters n_{max} and $n_{1/2}$ from the LINPACK benchmark. This results in the matrices shown in Figure 1 for the class A problem sizes and in Figure 2 for the class B results.

	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP
r_{max}	0.99										
n_{max}	0.87	0.81									
$n_{1/2}$	0.60	0.50	0.66								
EP	0.97	0.97	0.82	0.53							
MG	0.90	0.94	0.75	0.36	0.92						
CG	0.58	0.62	0.48	0.20	0.63	0.73					
FT	0.92	0.95	0.66	0.47	0.93	0.98	0.77				
IS	0.60	0.63	0.52	0.11	0.65	0.75	0.99	0.78			
LU	0.93	0.95	0.93	0.55	0.93	0.98	0.77	0.98	0.81		
SP	0.94	0.96	0.86	0.41	0.95	0.96	0.78	0.98	0.80	0.99	
BT	0.88	0.92	0.58	0.42	0.86	0.98	0.67	0.98	0.68	0.96	0.96

Figure 1. The correlation matrix of the NAS PBs of class A and the LINPACK benchmark results.

You can see that the benchmark results and the peak performance are highly correlated in almost all cases except for CG and IS. The correlations between benchmark results and the parameters n_{max} and $n_{1/2}$ are on the average much smaller. Only n_{max} shows bigger correlations to some of the benchmarks. We found later on during our studies no evidence that these two parameters can be used to explain or determine benchmark results and did not include them in the later analyses.

The reason for the high correlations between benchmarks is the simple fact, that published benchmark results always improve with increasing system size. This is not very surprising as other results would not be published by vendors. These big correlations are the reason for problems in the factor analyses and their interpretation. They lead to a very dominating single factor which tends to hide all other effects. Therefore we

r_{max}	0.99										
n_{max}	0.87	0.82									
$n_{1/2}$	0.75	0.71	0.76								
EP	0.97	0.97	0.83	0.82							
MG	0.91	0.93	0.77	0.65	0.91						
CG	0.57	0.55	0.41	0.18	0.57	0.69					
FT	0.85	0.90	0.53	0.64	0.87	0.95	0.77				
IS	0.65	0.65	0.52	0.25	0.64	0.78	0.99	0.84			
LU	0.94	0.94	0.90	0.65	0.90	0.99	0.75	0.96	0.84		
SP	0.96	0.96	0.91	0.76	0.95	0.99	0.70	0.96	0.79	0.99	
BT	0.82	0.99	0.94	0.78	0.99	0.98	0.63	0.95	0.72	0.98	0.99
	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP

Figure 2. The correlation matrix of the NAS PBs of class B and the LINPACK benchmark results.

also used the correlation matrix between the efficiencies of the benchmarks with respect to the peak performance to eliminate the strong effect of the overall correlation to peak performance (Figure 3 and Figure 4). On the average the individual correlations are now smaller but in general still high. The peak performance shows no big correlations (neither positive nor negative!) to any of the efficiencies of the benchmarks. The correlations between the parameter $n_{1/2}$ and the efficiencies are now negative and stronger compared to previous cases but still not as high as the benchmark correlations. This means that no general conclusions about the efficiencies of benchmarks can be made from the LINPACK parameters n_{max} and $n_{1/2}$.

r_{max}	-0.06										
n_{max}	0.87	-0.17									
$n_{1/2}$	0.60	-0.55	0.66								
EP	-0.14	0.46	-0.15	-0.29							
MG	-0.18	0.78	-0.22	-0.46	0.61						
CG	-0.22	0.70	-0.32	-0.59	0.42	0.78					
FT	-0.16	0.74	-0.21	-0.51	0.57	0.89	0.88				
IS	-0.25	0.79	-0.27	-0.51	0.55	0.85	0.93	0.93			
LU	-0.22	0.79	-0.21	-0.58	0.29	0.86	0.86	0.83	0.84		
SP	-0.19	0.78	-0.18	-0.61	0.35	0.87	0.91	0.88	0.91	0.95	
BT	0.06	0.81	-0.07	-0.38	0.44	0.85	0.61	0.73	0.77	0.83	0.75
	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP

Figure 3. The correlation matrix of the efficiencies of the NAS PBs of class A and the LINPACK benchmark results.

The factor analyses of the benchmarks

Factor analyses can be used as an explorative method to get an overview on the structure of the given data, but cannot be used for testing or proofing any hypothesis. Therefore much care must be taken in interpreting the results.

results of the class A problem size. The result is checked by looking at linear regressions between different NAS PBs.

Amdahl's Law [4] gives a very simple model for the performance of a parallel system for different numbers of processors. We fitted the measured data to Amdahl's Law to see whether this is possible and whether sufficient statistical space for including additional parameters remains in this model. All analyses were done using the SAS statistical software package. The data used for the analyses in this paper is as of October 1994

The correlation matrices

As starting point for the factor analyses we had to calculate the correlation matrix. We used the NAS PB results of the class A benchmarks and the LINPACK results r_{max} from table 3 of the LINPACK report [3] for unlimited problem sizes. We included also the peak performance r_{peak}, and the parameters n_{max} and $n_{1/2}$ from the LINPACK benchmark. This results in the matrices shown in Figure 1 for the class A problem sizes and in Figure 2 for the class B results.

r_{max}	0.99										
n_{max}	0.87	0.81									
$n_{1/2}$	0.60	0.50	0.66								
EP	0.97	0.97	0.82	0.53							
MG	0.90	0.94	0.75	0.36	0.92						
CG	0.58	0.62	0.48	0.20	0.63	0.73					
FT	0.92	0.95	0.66	0.47	0.93	0.98	0.77				
IS	0.60	0.63	0.52	0.11	0.65	0.75	0.99	0.78			
LU	0.93	0.95	0.93	0.55	0.93	0.98	0.77	0.98	0.81		
SP	0.94	0.96	0.86	0.41	0.95	0.96	0.78	0.98	0.80	0.99	
BT	0.88	0.92	0.58	0.42	0.86	0.98	0.67	0.98	0.68	0.96	0.96
	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP

Figure 1. The correlation matrix of the NAS PBs of class A and the LINPACK benchmark results.

You can see that the benchmark results and the peak performance are highly correlated in almost all cases except for CG and IS. The correlations between benchmark results and the parameters n_{max} and $n_{1/2}$ are on the average much smaller. Only n_{max} shows bigger correlations to some of the benchmarks. We found later on during our studies no evidence that these two parameters can be used to explain or determine benchmark results and did not include them in the later analyses.

The reason for the high correlations between benchmarks is the simple fact, that published benchmark results always improve with increasing system size. This is not very surprising as other results would not be published by vendors. These big correlations are the reason for problems in the factor analyses and their interpretation. They lead to a very dominating single factor which tends to hide all other effects. Therefore we

6). We get eigenvalues of 6.2, 0.7 and 0.05 by extracting 3 factors. Factor 2 has high loadings from CG and IS and factor 1 from all other benchmarks. For factor 3 a safe interpretation cannot be made. None of the experiments with factor analyses ever showed an indication for more than four meaningful and independent factors in the group of nine benchmarks and the peak performance.

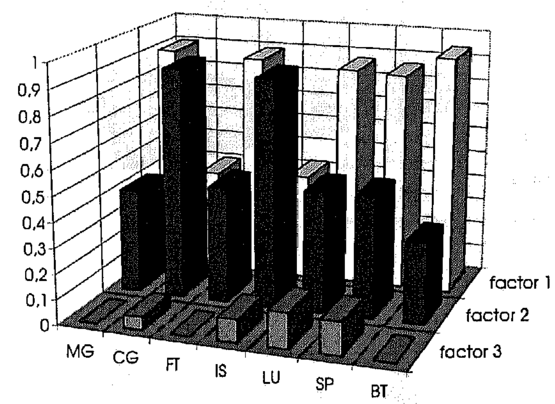

Figure 6. The loading of the 3 factors of the factor analysis for the subset of 7 NAS PBs only.

In addition to the performances we now analyze the correlations of the efficiencies in the same way. The result of the factor analysis of all benchmarks is shown in Figure 7. The first three eigenvalues of the correlation matrix are now 7.1, 0.9 and 0.5. Factor 3 shows a high loading from the efficiency of EP. Factor 1 and factor 2 show some different mix of the other benchmarks. Factor 1 has again high loadings from CG and IS which are not present in the other factors.

In Figure 8 the seven NAS PBs without EP show high loadings of CG and IS and in addition also of FT in factor 1. Factor 2 shows high loadings from MG and BT and factor 3 from LU and SP. Taking into consideration that FT was not loading high together with CG and IS in the case of the benchmark data we therefore summarize the factor analysis as follows:

- All benchmarks are strongly correlated with the peak performance. The different factor analyses indicate at the most four independent factors.
- r_{max} from the LINPACK benchmark, EP and the peak performance are highly correlated and as a group form one factor of the analyses.
- CG and IS as a group always form a second factor in the analyses.
- The remaining five NAS Parallel Benchmarks can be arranged in the two groups (LU and SP) and (MG, FT and BT). But the statistical evidence for this splitting is not as clear as for the other groups.

A common rule of experience demands at least 50 observations for applying a factor analysis at all. As the number of complete sets of mea-

	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP
r_{max}	-0.09										
n_{max}	0.87	-0.22									
$n_{1/2}$	0.75	-0.39	0.76								
EP	-0.22	0.15	-0.21	0.04							
MG	-0.28	0.68	-0.43	-0.42	0.13						
CG	-0.25	0.70	-0.42	-0.49	0.31	0.89					
FT	-0.11	0.48	-0.26	-0.31	0.36	0.74	0.82				
IS	-0.30	0.73	-0.44	-0.52	0.41	0.87	0.99	0.82			
LU	-0.35	0.70	-0.43	-0.53	-0.03	0.95	0.90	0.67	0.88		
SP	-0.32	0.70	-0.48	-0.56	0.10	0.92	0.93	0.66	0.89	0.94	
BT	-0.07	0.63	-0.52	-0.57	0.51	0.86	0.89	0.69	0.90	0.93	0.96

Figure 4. The correlation matrix of the efficiencies of the NAS PBs of class B and the LINPACK benchmark results.

First we applied a factor analysis to the correlation matrix of the benchmark results shown in Figure 1. The LINPACK parameters n_{max} and $n_{1/2}$ always came out as individual factors and thus gave us no additional information, so we did not include these two variables in the factor analyses any more. This first factor analysis gives a very dominant factor (with an eigenvalue of 8.8), which is related to the overall increase of benchmark performance with respect to the increase in peak performance. The next eigenvalues of the correlation matrix are 0.9 and 0.2 and thus for a rigid interpretation already quite small. Extracting these 3 factors and looking on the loading of their components after rotation (Figure 5) you can see that factor 2 has high loadings from CG and IS and factor 1 from EP, r_{max} and r_{peak}, while factor 3 contains medium loadings from all other benchmarks.

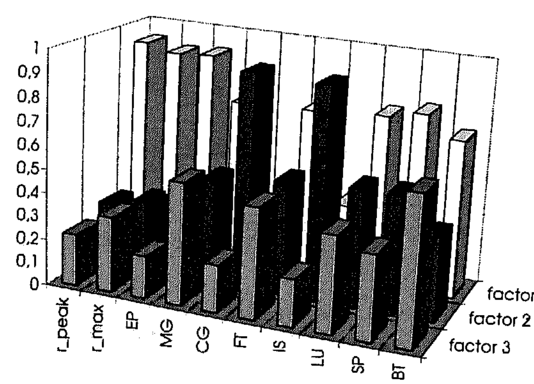

Figure 5. The loading of the 3 factors of the factor analysis.

To get more informations about the seven NAS PB not including EP we performed a second factor analysis on this group of benchmarks (Figure

6). We get eigenvalues of 6.2, 0.7 and 0.05 by extracting 3 factors. Factor 2 has high loadings from CG and IS and factor 1 from all other benchmarks. For factor 3 a safe interpretation cannot be made. None of the experiments with factor analyses ever showed an indication for more than four meaningful and independent factors in the group of nine benchmarks and the peak performance.

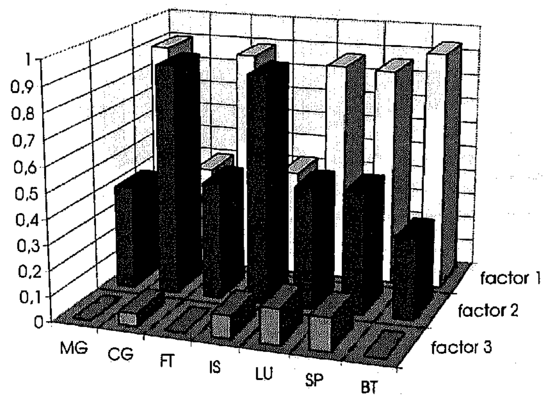

Figure 6. The loading of the 3 factors of the factor analysis for the subset of 7 NAS PBs only.

In addition to the performances we now analyze the correlations of the efficiencies in the same way. The result of the factor analysis of all benchmarks is shown in Figure 7. The first three eigenvalues of the correlation matrix are now 7.1, 0.9 and 0.5. Factor 3 shows a high loading from the efficiency of EP. Factor 1 and factor 2 show some different mix of the other benchmarks. Factor 1 has again high loadings from CG and IS which are not present in the other factors.

In Figure 8 the seven NAS PBs without EP show high loadings of CG and IS and in addition also of FT in factor 1. Factor 2 shows high loadings from MG and BT and factor 3 from LU and SP. Taking into consideration that FT was not loading high together with CG and IS in the case of the benchmark data we therefore summarize the factor analysis as follows:

- All benchmarks are strongly correlated with the peak performance. The different factor analyses indicate at the most four independent factors.
- r_{max} from the LINPACK benchmark, EP and the peak performance are highly correlated and as a group form one factor of the analyses.
- CG and IS as a group always form a second factor in the analyses.
- The remaining five NAS Parallel Benchmarks can be arranged in the two groups (LU and SP) and (MG, FT and BT). But the statistical evidence for this splitting is not as clear as for the other groups.

A common rule of experience demands at least 50 observations for applying a factor analysis at all. As the number of complete sets of mea-

	r_{peak}	r_{max}	n_{max}	$n_{1/2}$	EP	MG	CG	FT	IS	LU	SP
r_{max}	-0.09										
n_{max}	0.87	-0.22									
$n_{1/2}$	0.75	-0.39	0.76								
EP	-0.22	0.15	-0.21	0.04							
MG	-0.28	0.68	-0.43	-0.42	0.13						
CG	-0.25	0.70	-0.42	-0.49	0.31	0.89					
FT	-0.11	0.48	-0.26	-0.31	0.36	0.74	0.82				
IS	-0.30	0.73	-0.44	-0.52	0.41	0.87	0.99	0.82			
LU	-0.35	0.70	-0.43	-0.53	-0.03	0.95	0.90	0.67	0.88		
SP	-0.32	0.70	-0.48	-0.56	0.10	0.92	0.93	0.66	0.89	0.94	
BT	-0.07	0.63	-0.52	-0.57	0.51	0.86	0.89	0.69	0.90	0.93	0.96

Figure 4. The correlation matrix of the efficiencies of the NAS PBs of class B and the LINPACK benchmark results.

First we applied a factor analysis to the correlation matrix of the benchmark results shown in Figure 1. The LINPACK parameters n_{max} and $n_{1/2}$ always came out as individual factors and thus gave us no additional information, so we did not include these two variables in the factor analyses any more. This first factor analysis gives a very dominant factor (with an eigenvalue of 8.8), which is related to the overall increase of benchmark performance with respect to the increase in peak performance. The next eigenvalues of the correlation matrix are 0.9 and 0.2 and thus for a rigid interpretation already quite small. Extracting these 3 factors and looking on the loading of their components after rotation (Figure 5) you can see that factor 2 has high loadings from CG and IS and factor 1 from EP, r_{max} and r_{peak}, while factor 3 contains medium loadings from all other benchmarks.

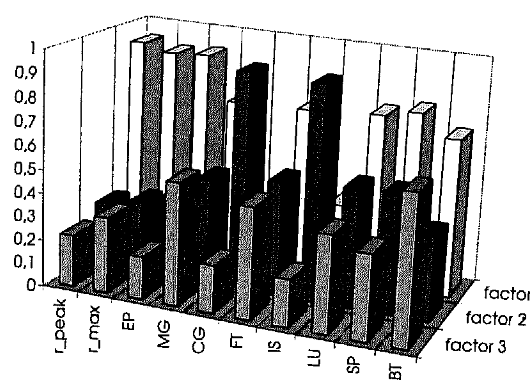

Figure 5. The loading of the 3 factors of the factor analysis.

To get more informations about the seven NAS PB not including EP we performed a second factor analysis on this group of benchmarks (Figure yses in the section above, we calculated the linear regressions between different pairs of benchmarks. As we never saw a statistical significance for an intercept term, we excluded it from the fit. Thus each regression is characterized by two parameters:

– β: the slope of the regression line.
– R^2: the portion of the variance σ explained by the regression.

In the lower left of Figure 9 we show the slope β, in the upper right we show R^2. All R^2 values are quite high. The regressions are better for pairs of benchmarks from within one of the identified groups than otherwise.

ure 9. The slope of regression lines β d the R^2 values for pairwise regressions. values are shown in lower left corner d R^2 values in the per right. r_{peak} and ax are measured Mflop/s while the AS PBs are given in AS PB units.

$\beta\backslash R^2$	r_{peak}	r_{max}	EP	MG	CG	FT	IS	LU	SP	BT
r_{peak}		0.998	0.990	0.950	0.426	0.936	0.464	0.884	0.923	0.990
r_{max}	0.66333		0.993	0.963	0.458	0.951	0.497	0.903	0.941	0.996
EP	0.00152	0.00230		0.961	0.501	0.960	0.536	0.908	0.950	0.989
MG	0.00063	0.00096	0.417		0.589	0.979	0.641	0.981	0.979	0.974
CG	0.00018	0.00027	0.124	0.32		0.641	0.987	0.677	0.679	0.484
FT	0.00064	0.00098	0.426	1.01	1.98		0.684	0.972	0.994	0.967
IS	0.00023	0.00036	0.163	0.42	1.26	0.42		0.734	0.721	0.527
LU	0.00036	0.00055	0.240	0.59	1.18	0.57	0.97		0.978	0.926
SP	0.00063	0.00095	0.414	0.99	1.99	0.97	1.62	1.67		0.958
BT	0.00117	0.00177	0.764	1.78	3.04	1.74	2.51	2.93	1.77	

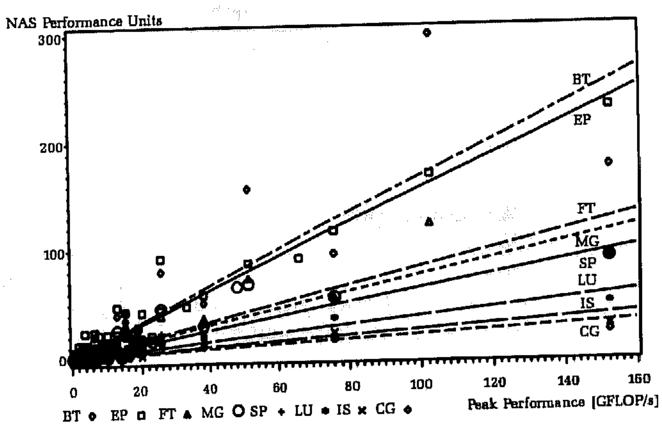

igure 10. Linear egression of all NAS PBs versus the peak performance.

For a possible interpretation of the result of the factor analyses, we then plotted all NAS PB results over the peak performance (Figure 10) and

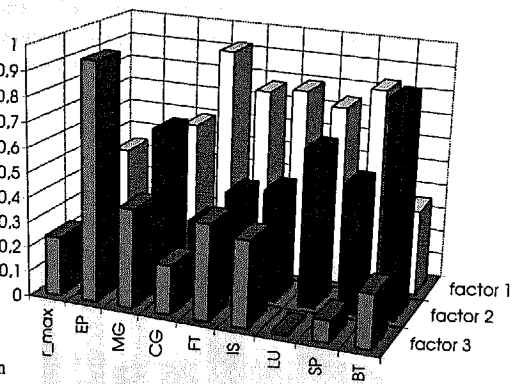

Figure 7. The loading of the 3 factors of the factor analysis based on efficiencies.

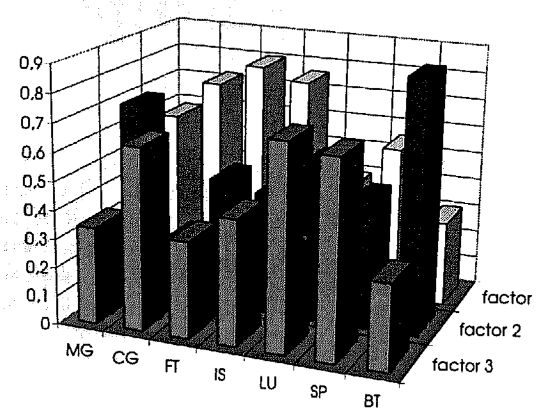

Figure 8. The loading of the 3 factors of the factor analysis for the subset of 7 NAS PBs only based on efficiencies.

surements for all benchmarks (about 30) is quite low compared to this, we calculated each coefficient of the correlation matrices by using also incomplete observations. This gives on the average 59 observations per element. But now the correlation matrices can have negative eigenvalues which might make a factor analysis meaningless. In our case the absolute values of the negative eigenvalues are very small and the dominating factors and their components are quite similar to the ones obtained by using complete observations only. So we take this as an additional confirmation of our analyses. For the class B problem size on the average only 30 observations are available for the correlation matrices. Due to this small statistical basis we do not report results of the factor analyses based on class B results. But at a first look they seem to be similar.

For a first check of the groups of benchmarks identified by the factor analyses in the section above, we calculated the linear regressions between different pairs of benchmarks. As we never saw a statistical significance for an intercept term, we excluded it from the fit. Thus each regression is characterized by two parameters:

- β: the slope of the regression line.
- R^2: the portion of the variance σ explained by the regression.

In the lower left of Figure 9 we show the slope β, in the upper right we show R^2. All R^2 values are quite high. The regressions are better for pairs of benchmarks from within one of the identified groups than otherwise.

$\beta\backslash R^2$	r_{peak}	r_{max}	EP	MG	CG	FT	IS	LU	SP	BT
r_{peak}		0.998	0.990	0.950	0.426	0.936	0.464	0.884	0.923	0.990
r_{max}	0.66333		0.993	0.963	0.458	0.951	0.497	0.903	0.941	0.996
EP	0.00152	0.00230		0.961	0.501	0.960	0.536	0.908	0.950	0.989
MG	0.00063	0.00096	0.417		0.589	0.979	0.641	0.981	0.979	0.974
CG	0.00018	0.00027	0.124	0.32		0.641	0.987	0.677	0.679	0.484
FT	0.00064	0.00098	0.426	1.01	1.98		0.684	0.972	0.994	0.967
IS	0.00023	0.00036	0.163	0.42	1.26	0.42		0.734	0.721	0.527
LU	0.00036	0.00055	0.240	0.59	1.18	0.57	0.97		0.978	0.926
SP	0.00063	0.00095	0.414	0.99	1.99	0.97	1.62	1.67		0.958
BT	0.00117	0.00177	0.764	1.78	3.04	1.74	2.51	2.93	1.77	

Figure 9. The slope of the regression lines β and the R^2 values for all pairwise regressions. β values are shown in the lower left corner and R^2 values in the upper right. r_{peak} and r_{max} are measured in Mflop/s while the NAS PBs are given in NAS PB units.

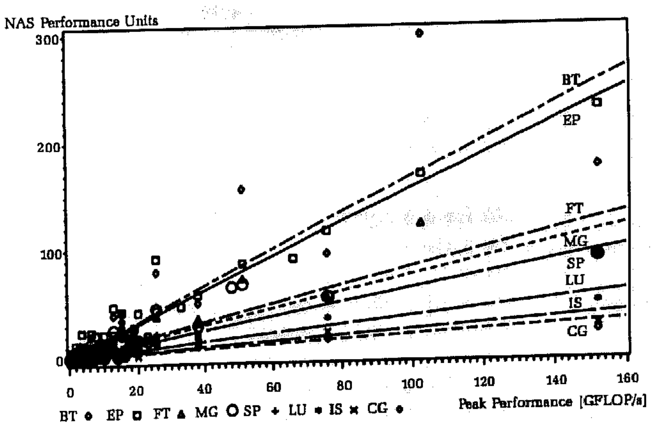

Figure 10. Linear regression of all NAS PBs versus the peak performance.

For a possible interpretation of the result of the factor analyses, we then plotted all NAS PB results over the peak performance (Figure 10) and

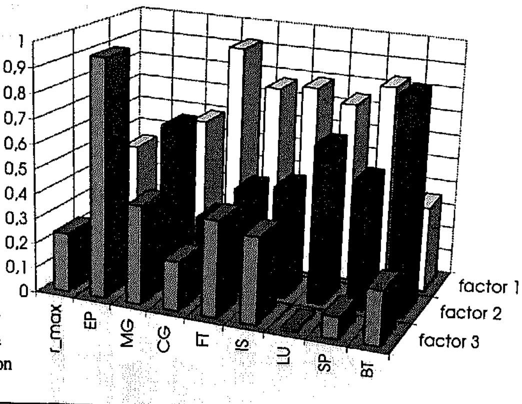

Figure 7. The loading of the 3 factors of the factor analysis based on efficiencies.

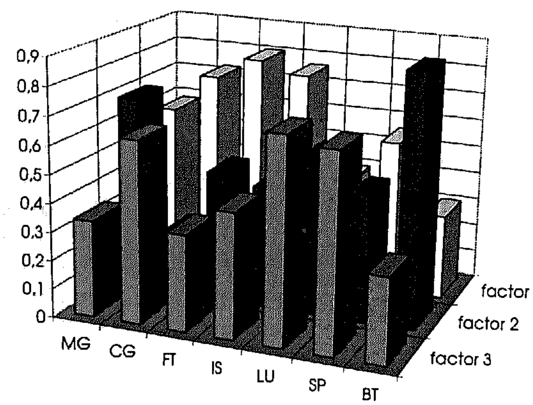

Figure 8. The loading of the 3 factors of the factor analysis for the subset of 7 NAS PBs only based on efficiencies.

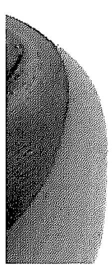

surements for all benchmarks (about 30) is quite low compared to this, we calculated each coefficient of the correlation matrices by using also incomplete observations. This gives on the average 59 observations per element. But now the correlation matrices can have negative eigenvalues which might make a factor analysis meaningless. In our case the absolute values of the negative eigenvalues are very small and the dominating factors and their components are quite similar to the ones obtained by using complete observations only. So we take this as an additional confirmation of our analyses. For the class B problem size on the average only 30 observations are available for the correlation matrices. Due to this small statistical basis we do not report results of the factor analyses based on class B results. But at a first look they seem to be similar.

For a first check of the groups of benchmarks identified by the factor anal-

r_1 of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	0.00409	0.00103	0.00217	0.00185	0.00085	0.00121	0.00084	0.00151
CM5	0.18736	0.03593	0.02024	0.09819	0.00778	0.02894	0.06697	0.10990
CM5E	0.35747	0.19657	0.03644	0.17422	0.06013	0.08540	0.09302	0.17805
KSR1	0.07652	0.03702				0.04660	0.03918	0.05632
Meiko CS2	0.20828	0.20783		0.16815				
nCube 2s	0.02351	0.00897	0.00633	0.00677	0.00764	0.00403	0.00538	0.00981
SGI PowChal	0.51472		0.48834	0.51727		0.41571	0.55089	0.57377
IBM SP1		0.15804	0.06619	0.08340	0.09351	0.13906	0.13248	0.21735
IBM SP2	0.35536	0.43788	0.38369	0.21532	0.25259	0.36609	0.33067	0.41855
SPP1000	0.32955	0.10930	0.06169	0.16566	0.14220	0.15734	0.18916	0.29769
Cray T3D	0.22689	0.13125	0.04925	0.15273	0.08238	0.10443	0.13640	0.20387
VPP500	2.78729	3.95379	2.25813	2.66994	4.99801		2.54733	5.07530
Paragon XP	0.19190	0.05103	0.10621	0.07056		0.03230	0.04044	0.06887
Y-MP C90	2.72625	3.03040	3.45228	3.08355	3.58960	2.21656	2.41429	2.16728
Y-MPel		0.30116	0.28088	0.31413	0.27584	0.22234	0.27992	0.23899

gure 11. Single ocessor performances in NAS PB units tained by a fit Amdahl's Law all NAS PBs for ass A problem size. issing values indicate easurements for only vo or less system zes.

α of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	0.99988	0.99984	0.99937	0.99976	0.99939	0.99898	0.99934	0.99978
CM5	0.99988	0.99780	0.99727	0.98578	0.99968	0.99281	0.99089	0.99122
CM5E	0.99932	0.99592	0.99821	0.99039	0.99910	0.99092	0.99563	0.99697
KSR1	1.00008	0.99856				0.98677	0.99421	0.99754
Meiko CS2	0.99741	0.99439		0.98897				
nCube 2s	1.00000	0.99984	0.99877	1.00001	0.99985	0.99935	0.99959	0.99973
SGI PowChal	0.99930		0.91776	0.86025		0.98009	0.98286	0.99578
IBM SP1		0.99498	0.98871	0.99700	0.99021	0.98916	0.98699	0.99262
IBM SP2	1.00005	0.99583	0.95535	0.99733	0.99484	0.98805	0.99102	0.99445
SPP1000	0.99894	0.96935	1.00027	0.97542	0.93409	0.96209	0.96072	0.98246
Cray T3D	0.99999	0.99950	0.99858	0.99926	0.99788	0.99879	0.99941	0.99980
VPP500	0.99901	0.97157	0.90730	0.99368	0.68633		0.94725	0.99857
Paragon XP	0.99985	0.99704	0.97843	0.99517		0.99661	0.99551	0.99648
Y-MP C90	0.99873	0.92711	0.96165	0.96262	0.97633	0.94180	0.99533	0.98383
Y-MPel		0.79131	0.87243	0.91428	0.91003	0.87466	0.82094	0.88249

Figure 12. Parallelization ratios α obtained by a fit of Amdahl's Law to all NAS PBs for class A problem size. Missing values indicate measurements for only two or less system sizes.

this case Amdahl's Law is too limited and cannot be extrapolated to unlimited processor numbers. So transformation of equation (2) fails for these cases.

The resulting parameters shown in Figures 11–14 give a good characterization and overview on the different systems and on the implementations of the benchmarks. For instance, is it quite easy to see extraordinarily good or bad implementations and results.

For the class B benchmark sizes we fitted the parameters shown in Figures 15–18.

In most cases Amdahl's Law fits very well to the data, giving small error bounds for any prediction typically in the range of a few percent. Thus calculated a linear regression over all results for different systems for each single benchmark. From top to bottom we found the ordering BT, EP, FT, MG, SP, LU, CG and IS. BT shows up higher than EP only because of the very well tuned results for the Fujitsu VPP500 system.So the different groups of benchmarks appear next to each other, giving a first interpretation of the results of the factor analyses: The different groups of benchmarks have different characteristic ranges of efficiencies of their implementations.

Applying Amdahl's Law to the NAS PBs

One of the simplest known models for the performance of a problem of fixed size on a parallel system is Amdahl's Law [4]. It's basic assumption is a split of the computational work in a sequential and in a fully parallelizable part. It can be characterized by the following parametrizations:

- α Fraction of parallelizable work in the implementation of the code
- $1-\alpha$ Fraction of sequential work in the implementation of the code
- r_1 Performance running the code on a single processor given in units of the NAS PBs, which are dimensionless.
- N Number of processors used
- t_N Time for executing the program using N processors
- Sp_N Speedup of the program on N processors compared to one processor:

$$Sp_N = \frac{t_1}{t_N} = \frac{r_1 * N}{N - \alpha(N-1)} \tag{1}$$

A different parametrization can be obtained by introducing the asymptotic performance r_∞ achieved by using an infinite number of processors and the processor number $N_{1/2}$ needed for achieving half of r_∞ as follows:

$$r_\infty = \frac{r_1}{1-\alpha} \text{ and } N_{1/2} = \frac{\alpha}{1-\alpha} \tag{2}$$

This gives:

$$Sp_N = \frac{r_\infty}{1 + \frac{N_{1/2}}{N}}$$

We fitted Amdahl's Law to the NAS PB to look whether this simple model for performance is already able to explain the measured performances or whether there is statistical room for more sophisticated models. As we wanted to calculate error terms we did this only for systems for which performance data of at least three different system sizes are reported. We allowed also α values greater than one, which does not make sense in a rigid application of Amdahl's Law. The results are shown in Figures 11 and 12 for the parametrization given in equation (1). By applying the transformation of equation (2) these values can be transformed into Figures 13 and 14. Some of the systems show α values slightly greater than one. This can be seen as an indication of superlinear speedups. In

r_1 of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	0.00409	0.00103	0.00217	0.00185	0.00085	0.00121	0.00084	0.00151
CM5	0.18736	0.03593	0.02024	0.09819	0.00778	0.02894	0.06697	0.10990
CM5E	0.35747	0.19657	0.03644	0.17422	0.06013	0.08540	0.09302	0.17805
KSR1	0.07652	0.03702				0.04660	0.03918	0.05632
Meiko CS2	0.20828	0.20783		0.16815				
nCube 2s	0.02351	0.00897	0.00633	0.00677	0.00764	0.00403	0.00538	0.00981
SGI PowChal	0.51472		0.48834	0.51727		0.41571	0.55089	0.57377
IBM SP1		0.15804	0.06619	0.08340	0.09351	0.13906	0.13248	0.21735
IBM SP2	0.35536	0.43788	0.38369	0.21532	0.25259	0.36609	0.33067	0.41855
SPP1000	0.32955	0.10930	0.06169	0.16566	0.14220	0.15734	0.18916	0.29769
Cray T3D	0.22689	0.13125	0.04925	0.15273	0.08238	0.10443	0.13640	0.20387
VPP500	2.78729	3.95379	2.25813	2.66994	4.99801		2.54733	5.07530
Paragon XP	0.19190	0.05103	0.10621	0.07056		0.03230	0.04044	0.06887
Y-MP C90	2.72625	3.03040	3.45228	3.08355	3.58960	2.21656	2.41429	2.16728
Y-MPel		0.30116	0.28088	0.31413	0.27584	0.22234	0.27992	0.23899

Figure 11. Single processor performances r_1 in NAS PB units obtained by a fit of Amdahl's Law to all NAS PBs for class A problem size. Missing values indicate measurements for only two or less system sizes.

α of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	0.99988	0.99984	0.99937	0.99976	0.99939	0.99898	0.99934	0.99978
CM5	0.99988	0.99780	0.99727	0.98578	0.99968	0.99281	0.99089	0.99122
CM5E	0.99932	0.99592	0.99821	0.99039	0.99910	0.99092	0.99563	0.99697
KSR1	1.00008	0.99856				0.98677	0.99421	0.99754
Meiko CS2	0.99741	0.99439		0.98897				
nCube 2s	1.00000	0.99984	0.99877	1.00001	0.99985	0.99935	0.99959	0.99973
SGI PowChal	0.99930		0.91776	0.86025		0.98009	0.98286	0.99578
IBM SP1		0.99498	0.98871	0.99700	0.99021	0.98916	0.98699	0.99262
IBM SP2	1.00005	0.99583	0.95535	0.99733	0.99484	0.98805	0.99102	0.99445
SPP1000	0.99894	0.96935	1.00027	0.97542	0.93409	0.96209	0.96072	0.98246
Cray T3D	0.99999	0.99950	0.99858	0.99926	0.99788	0.99879	0.99941	0.99980
VPP500	0.99901	0.97157	0.90730	0.99368	0.68633		0.94725	0.99857
Paragon XP	0.99985	0.99704	0.97843	0.99517		0.99661	0.99551	0.99648
Y-MP C90	0.99873	0.92711	0.96165	0.96262	0.97633	0.94180	0.99533	0.98383
Y-MPel		0.79131	0.87243	0.91428	0.91003	0.87466	0.82094	0.88249

Figure 12. Parallelization ratios α obtained by a fit of Amdahl's Law to all NAS PBs for class A problem size. Missing values indicate measurements for only two or less system sizes.

this case Amdahl's Law is too limited and cannot be extrapolated to unlimited processor numbers. So transformation of equation (2) fails for these cases.

The resulting parameters shown in Figures 11–14 give a good characterization and overview on the different systems and on the implementations of the benchmarks. For instance, is it quite easy to see extraordinarily good or bad implementations and results.

For the class B benchmark sizes we fitted the parameters shown in Figures 15–18.

In most cases Amdahl's Law fits very well to the data, giving small error bounds for any prediction typically in the range of a few percent. Thus calculated a linear regression over all results for different systems for each single benchmark. From top to bottom we found the ordering BT, EP, FT, MG, SP, LU, CG and IS. BT shows up higher than EP only because of the very well tuned results for the Fujitsu VPP500 system.So the different groups of benchmarks appear next to each other, giving a first interpretation of the results of the factor analyses: The different groups of benchmarks have different characteristic ranges of efficiencies of their implementations.

Applying Amdahl's Law to the NAS PBs

One of the simplest known models for the performance of a problem of fixed size on a parallel system is Amdahl's Law [4]. It's basic assumption is a split of the computational work in a sequential and in a fully parallelizable part. It can be characterized by the following parametrizations:

- α Fraction of parallelizable work in the implementation of the code
- $1-\alpha$ Fraction of sequential work in the implementation of the code
- r_1 Performance running the code on a single processor given in units of the NAS PBs, which are dimensionless.
- N Number of processors used
- t_N Time for executing the program using N processors
- Sp_N Speedup of the program on N processors compared to one processor:

$$Sp_N = \frac{t_1}{t_N} = \frac{r_1 * N}{N - \alpha(N-1)} \quad (1)$$

A different parametrization can be obtained by introducing the asymptotic performance r_∞ achieved by using an infinite number of processors and the processor number $N_{1/2}$ needed for achieving half of r_∞ as follows:

$$r_\infty = \frac{r_1}{1-\alpha} \text{ and } N_{1/2} = \frac{\alpha}{1-\alpha} \quad (2)$$

This gives:

$$Sp_N = \frac{r_\infty}{1 + \frac{N_{1/2}}{N}}$$

We fitted Amdahl's Law to the NAS PB to look whether this simple model for performance is already able to explain the measured performances or whether there is statistical room for more sophisticated models. As we wanted to calculate error terms we did this only for systems for which performance data of at least three different system sizes are reported. We allowed also α values greater than one, which does not make sense in a rigid application of Amdahl's Law. The results are shown in Figures 11 and 12 for the parametrization given in equation (1). By applying the transformation of equation (2) these values can be transformed into Figures 13 and 14. Some of the systems show α values slightly greater than one. This can be seen as an indication of superlinear speedups. In

r_1 of Class B	EP	MG	CG	FT	IS	LU	SP	BT
CM5E	0.12151	0.06228	0.00886	0.06382	0.01313	0.06463	0.02204	0.04965
Meiko CS2	0.07882	0.07026						
nCube 2s	0.00858				0.00203			
SGI PowChal	0.18901			0.19526		0.17161	0.21789	0.22499
IBM SP1		0.05711	0.01113	0.02773	0.01875	0.07758	0.05184	0.08320
IBM SP2	0.14902	0.15836	0.07889	0.07519	0.06636	0.17667	0.12340	0.16243
Cray T3D	0.08263	0.04764	0.01127	0.04742	0.01757	0.04155	0.04172	0.07342
VPP500	1.03391	1.39342	0.63652	1.04332			0.86942	1.99875
Paragon XP	0.06947	0.01743	0.00920		0.01562	0.01662	0.01278	0.02353
Y-MP C90	1.00674	1.13505	0.99427	2.65917	1.00900	1.03262	1.01801	1.02131

Figure 15. Single processor performances r_1 in NAS PB units obtained by a fit of Amdahl's Law to all NAS PBs for class B problem size. Missing values indicate measurements for only two or less system sizes.

α of Class B	EP	MG	CG	FT	IS	LU	SP	BT
CM5E	1.00020	0.99667	1.00100	0.99080	0.99906	0.97627	0.99774	0.99783
Meiko CS2	0.99885	0.99559						
nCube 2s	1.00000				0.99996			
SGI PowChal	0.99947			0.89185		0.97924	0.96530	0.98557
IBM SP1		0.99459	0.99688	0.99936	1.00000	0.99042	0.99476	0.99746
IBM SP2	1.00009	0.99674	0.99015	0.99919	0.99754	0.99415	0.99671	0.99809
Cray T3D	1.00000	0.99953	0.99945	0.99957	0.99932	0.99964	0.99949	0.99978
VPP500	0.99955	0.97301	0.95304	0.99777			0.97478	0.99961
Paragon XP	0.99998	0.99771	0.99841		0.99646	0.99745	0.99834	0.99829
Y-MP C90	0.99828	0.93933	0.97524	0.88526	0.98811	0.97925	0.94467	0.98351

Figure 16. Parallelization ratios α obtained by a fit of Amdahl's Law to all NAS PBs for class B problem size. Missing values indicate measurements for only two or less system sizes.

r_∞ of Class B	EP	MG	CG	FT	IS	LU	SP	BT
CM5E	*	18.7	*	6.9	14.0	2.7	9.6	22.9
Meiko CS2	66.5	16.2						
nCube 2s	*				50.75			
SGI PowChal	356.6			1.8		8.3	6.3	15.6
IBM SP1		10.6	3.6	43.3	*	8.1	9.9	32.8
IBM SP2	*	48.6	8.0	92.9	27.0	30.2	37.5	85.0
Cray T3D	*	101.4	20.5	110.3	25.8	115.4	81.8	333.7
VPP500	2297.6	51.7	13.6	505.4			34.5	5125.0
Paragon XP	3473.5	6.8	5.8		4.4	6.5	7.7	13.8
Y-MP C90	585.3	18.7	40.2	23.1	84.9	49.1	18.4	61.9

Figure 17. Asymptotic performances r_∞ in NAS PB units obtained by a fit of Amdahl's Law to all NAS PBs for class B problem size. Missing values indicate measurements for only two or less system sizes. "*" denote entries for which $\alpha \geq 1$.

Conclusions

Applying factor analyses to the NAS PBs we found that at the most four factors can be extracted for which a meaningful interpretation is possible. They add up to more than 95% of the total variance of the input data. Hence four benchmarks are sufficient to characterize the overall NAS Parallel Benchmark performances. Looking at the individual factors resulting from the analyses the results can be summarized as follows:

- All benchmarks are strongly correlated with the peak performance.
- r_{max} from the LINPACK benchmark, EP and the peak performance

r_∞ of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	34.1	6.4	3.4	0.8	1.4	1.2	1.3	6.9
CM5	1561.3	16.3	6.8	6.9	24.3	4.0	7.4	12.5
CM5E	526.0	48.2	20.4	18.1	66.8	9.4	21.3	58.8
KSR1	*	25.7				3.5	6.8	22.9
Meiko CS2	80.4	37.0		15.2				
nCube 2s	*	1.6	5.1	*	50.9	6.2	13.1	36.3
SGI PowChal	735.3		5.5	3.7		20.9	32.1	136.0
IBM SP1		31.5	5.9	22.8	9.6	12.8	10.2	29.3
IBM SP2	*	105.0	8.6	80.6	49.0	28.1	36.8	75.4
SPP1000	310.9	3.6	*	6.7	2.2	4.2	4.8	17.0
Cray T3D	2268.9	262.5	34.7	205.8	38.9	86.3	231.2	1019.4
VPP500	281.5	139.1	24.4	422.5	15.9		48.3	3549.2
Paragon XP	1279.3	17.2	4.9	14.6		9.5	9.0	19.6
Y-MP C90	2145.6	41.6	90.0	84.5	151.7	38.1	517.0	134.0
Y-MPel		1.4	2.2	3.7	3.1	1.8	1.6	2.0

Figure 13. Asymptotic performances r_∞ in NAS PB units obtained by a fit of Amdahl's Law to all NAS PBs. "*" denote entries for which $\alpha \geq 1$.

$N_{1/2}$ of Class A	EP	MG	CG	FT	IS	LU	SP	BT
CM2	8332.3	6249.0	1586.3	416.6	1683.3	979.4	1514.2	4544.6
CM5	8332.3	453.6	335.4	69.3	3124.0	138.1	108.8	112.9
CM5E	1469.6	244.1	557.7	103.1	1110.1	109.1	227.9	329.0
KSR1	*	693.4				74.6	171.7	405.5
Meiko CS2	385.1	177.3		89.7				
nCube 2s	*	178.2	821.0	*	6665.6	1537.5	2438.1	3702.7
SGI PowChal	1427.6		11.2	6.2		49.2	57.3	236.0
IBM SP1		198.2	87.6	332.3	101.1	91.3	75.9	134.0
IBM SP2	*	238.8	21.4	373.5	192.8	82.7	110.4	179.2
SPP1000	942.4	36.6	*	39.7	14.2	25.4	24.5	56.0
Cray T3D	99999.	1999.0	703.2	1350.4	470.7	825.4	1693.9	4999.0
VPP500	1009.1	34.2	9.8	157.2	2.2		18.0	698.3
Paragon XP	6665.7	336.8	45.4	206.0		294.0	221.7	283.1
Y-MP C90	77.9	12.7	25.1	25.1	41.2	16.2	213.1	60.8
Y-MPel		3.8	6.8	10.7	10.1	6.0	4.6	7.5

Figure 14. Processor number $N_{1/2}$ necessary for achieving half of the asymptotic performance r_∞ obtained by a fit of Amdahl's Law to all NAS PBs. "*" denote entries for which $\alpha \geq 1$.

no statistical space is left to include additional parameters in this model. This is not true in the case of the FT benchmark. We assume that this is due to the effect of parallelizing over a different number of dimensions of the FFT (1, 2 or 3 dimensions) for a different number of processors. If one fits only one unique Amdahl curve to the different domains of the implementation, this effect leads to a big error. Only by a closer look on the actual implementations, you would be able to decide if this explanation is true. Examples for the fitted curves for some systems and class A problem sizes are given in Figures 19–22. Here we also show as examples the error bounds for the fits of the BT and FT benchmarks. The errors for BT are typically quite small while the errors for FT are sometimes quite big.

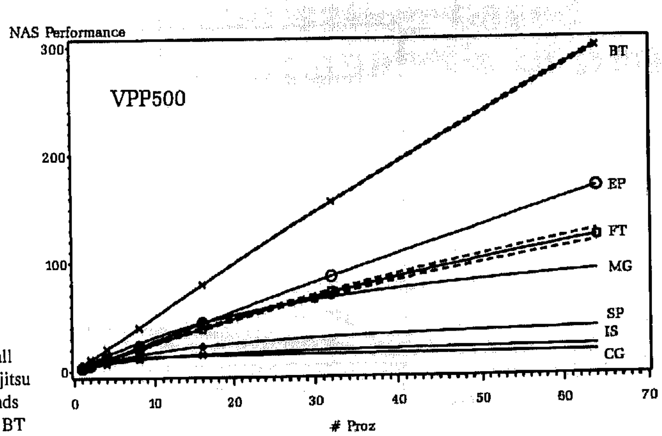

Figure 20. Fit of Amdahl's Law for all NAS PB for the Fujitsu VPP500. Error bounds are only shown for BT and FT.

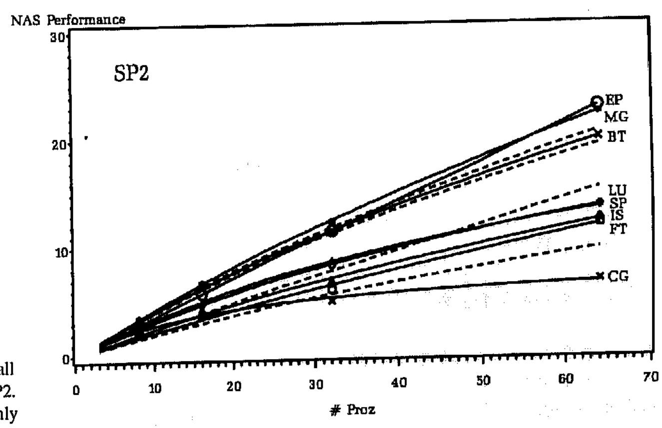

Figure 21. Fit of Amdahl's Law for all NAS PB for the SP2. Error bounds are only shown for BT and FT.

References

1
Bailey, D.H., J. Barton, T. Lasinski and H.D. Simon (eds.), *The NAS parallel benchmarks*, NASA Technical Memorandum **103863**, NASA Ames Research Center, Moffett Field CA 94035-1000, July 1993.

2
Bailey, D.H., E. Barszcz, L. Dagum and H.D. Simon, *NAS parallel benchmarks results 10-94*, NAS Technical Report **RNR-94-001**, NASA Ames Research Center, Moffett Field, CA 94035, October 1994.

are highly correlated and as a group form one factor of the analyses.
- CG and IS as a group always form a second factor in the analyses.
- The remaining five NAS Parallel Benchmarks can be arranged in the two groups (LU and SP) and (MG, FT and BT). But the statistical evidence for this splitting is not as clear as for the other groups.

We also used Amdahl's Law to fit the measured performances. In most cases Amdahl's Law fits very well to the data, giving small error bounds for any prediction. The resulting parallelization ratios, single processor performances, asymptotic performances and $N_{1/2}$ processor numbers give a good characterization and overview on the different systems and on the implementations of the benchmarks. For a more detailed analysis you would need access to the implementations of the codes used by the different vendors, but unfortunately those are proprietary to the vendors.

Figure 18. Processor number $N_{1/2}$ necessary for achieving half of the asymptotic performance r_∞ obtained by a fit of Amdahl's Law to all NAS PBs for class B problem size. Missing values indicate measurements for only two or less system sizes. "*" denote entries for which $\alpha \geq 1$.

$N_{1/2}$ of Class B	EP	MG	CG	FT	IS	LU	SP	BT
CM5E	*	299.3	*	107.7	1062.8	41.1	441.5	459.8
Meiko CS2	839.3	231.6						
nCube 2s	*				24991.0			
SGI PowChal	1885.8			8.2		47.2	27.8	68.3
IBM SP1		183.9	319.5	1561.5	*	103.4	189.8	392.7
IBM SP2	*	306.6	100.5	1233.6	405.5	169.9	303.0	522.6
Cray T3D	*	2126.7	1817.2	2324.6	1469.6	2776.8	1959.8	4544.5
VPP500	2221.2	36.1	20.3	483.3			38.7	2563.1
Paragon XP	49999.	435.7	627.9		281.5	391.2	601.4	583.8
Y-MP C90	580.4	15.5	39.4	7.7	83.1	47.2	17.1	59.6

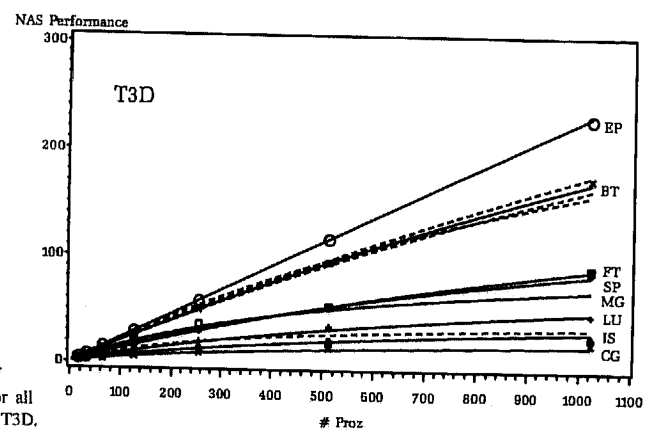

Figure 19. Fit of Amdahl's Law for all NAS PB for the T3D. Error bounds are only shown for BT and FT.

ution

Characterisation based bottleneck analysis of parallel systems

emerly, J. Papay,
udd

ity of Warwick,
Systems
Department
puter Science
AL Coventry, UK

mputer 62, XI–4

Bottleneck analysis plays an important role in the early design of parallel computers and programs. In this paper a methodology for bottleneck analysis based on an instruction level characterisation technique is presented. The methodology is based on the assumption that a bottleneck is caused by the slowest component of a computing system. These components are: memory (internal, external), processor (CPU, FPU), communication and I/O. Three metrics were used to identify bottlenecks in the system components. These are the B-ratio, the communication-computation ratio and the memory-processing ratio. These ratios are dimensionless and indicate the presence of a bottleneck when their values exceed unity. The methodology is illustrated and validated using a communication intensive linear solver algorithm (Gauss-Jordan elimination) which was implemented on a mesh connected distributed memory parallel computer (128 T800 Parsytec SuperCluster).

One of the main concerns of parallel computing is to port sequential programs efficiently knowing the resource limitations of the target machine such as processor, memory and communication network. In order to improve the performance of the parallel code bottleneck analysis is required. The identification of bottlenecks within parallel systems is an important aspect of hardware and software design. This process involves examining the system behavior under various load conditions. Bottlenecks can be defined in several ways as:

- The parts of the program that prevent achieving the optimal execution time.
- The parts of the system (either hardware or software) which consumes the maximum time or the slowest components of the system.

In this paper the second definition is used as the basis for the bottleneck analysis methodology which involves the following steps: predict the execution time components of a certain workload, identify the time component responsible for the bottleneck (the slowest part), analyze the component causing the bottleneck into its constituents and identify the sub-components causing the problem. Optimization of the software subroutines and/or hardware utilization causing the bottleneck can improve the system performance. This operation can be iterated until no further optimization is possible. Potential sources of bottlenecks are summarized

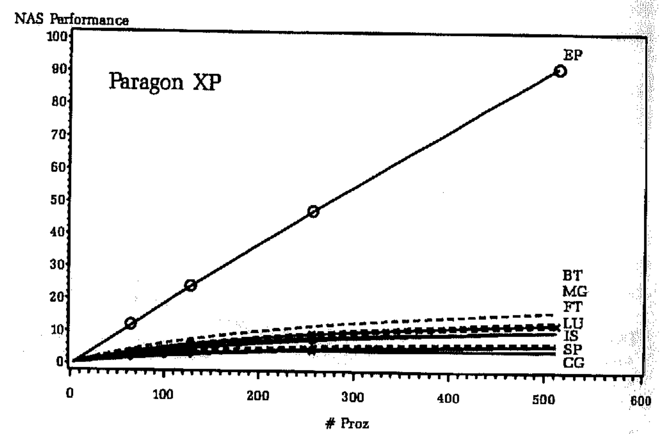

Figure 22. Fit of Amdahl's Law for all NAS PB for the Paragon XP running OSF1.2. Error bounds are only shown for BT and FT.

3
Dongarra, J.J., *Performance of Various Computers Using Standard Linear Equations Software*, Computer Science Department, University of Tennessee **CS-89-85**, 1994.

4
Amdahl, G.M., *Validity of the single processor approach to achieve large scale computing capabilities*, Proc. AFIPS 1967 Spring Joint Comput. Conf. **30**, 1967, 483–485.